

FORM PTO-1390 (REV 11-98)	U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE	APPlicant'S DOCKET NUMBER 36-1385
TRANSMITTAL LETTER TO THE UNITED STATES DESIGNATED/ELECTED OFFICE (DO/EO/US) CONCERNING A FILING UNDER 35 U.S.C. 371		U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.5) 09/674720 (To Be Assigned)
INTERNATIONAL APPLICATION NO. PCT/GB99/01773	INTERNATIONAL FILING DATE 4 June 1999	PRIORITY DATE CLAIMED 5 June 1998
TITLE OF INVENTION COMMUNICATIONS NETWORK		
APPLICANT(S) FOR DO/EO/US BRISCOE et al.		

Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:

- ☒ This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
- ☐ This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
- ☒ This is an express request to begin national examination procedures (35 U.S.C. 371(f) at any time rather than delay examination until the expiration of the applicable time limit set in 35 U.S.C. 371(b) and PCT Articles 22 and 39(1).
- ☒ A proper Demand for International Preliminary Examination was made by the 19th month from the earliest claimed priority date.
- A copy of the International Application as filed (35 U.S.C. 371(c)(2)).
 - ☒ is transmitted herewith (required only if not transmitted by the International Bureau).
 - ☒ has been transmitted by the International Bureau.
 - ☐ is not required, as the application was filed in the United States Receiving Office (RO/US).
- ☐ A translation of the International Application into English (35 U.S.C. 371(c)(2)).
- ☒ Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3)).
 - ☒ are transmitted herewith (required only if not transmitted by the International Bureau).
 - ☒ have been transmitted by the International Bureau.
 - ☐ have not been made; however, the time limit for making such amendments has **NOT** expired.
 - ☐ have not been made and will not be made.
- ☐ A translation of the amendments to the claims under PCT Article 19 (U.S.C. 371(c)(3)).
- ☒ An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
- ☐ A translation of the annexes to the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11. To 16. Below concern document(s) or information included:

- ☐ An Information Disclosure Statement under 37 C.F.R. 1.97 and 1.98.
- ☒ An assignment document for recording. A separate cover sheet in compliance with 37 C.F.R. 3.28 and 3.31 is included.
- ☒ A FIRST preliminary amendment.
☐ A SECOND or SUBSEQUENT preliminary amendment.
- ☐ A substitute specification.
- ☐ A change of power of attorney and/or address letter.
- ☐ Other items or information.

U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.492(a)(1)-(5))

INTERNATIONAL APPLICATION NO.

ATTORNEY'S DOCKET NUMBER

09/16/74/20

PCT/GB99/01773

36-1385

17. ☒ The following fees are submitted:

BASIC NATIONAL FEE (37 C.F.R. 1.492(a)(1)-(5):

- Neither international preliminary examination fee (37 C.F.R. 1.482) nor international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO and International Search Report not prepared by the EPO or JPO\$1000.00
- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but International Search Report prepared by the EPO or JPO.....\$860.00
- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO\$710.00
- International preliminary examination fee paid to USPTO (37 C.F.R. 1.482) but all claims did not satisfy provisions of PCT Article 33(1)-(4).....\$690.00
- International preliminary examination fee paid to USPTO (37 C.F.R. 1.482) and all claims satisfied provisions of PCT Article 33(1)-(4).....\$100.00

ENTER APPROPRIATE BASIC FEE AMOUNT =

Surcharge of \$130.00 for furnishing the oath or declaration later than ☐ 20 ☐ 30 months from the earliest claimed priority date (37 C.F.R. 1.492(e)).

CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE
Total Claims	17	-20 = 0	X \$18.00
Independent Claims	5	-3 = 2	X \$80.00

MULTIPLE DEPENDENT CLAIMS(S) (if applicable)

\$270.00

TOTAL OF ABOVE CALCULATIONS =

Reduction by 1/2 for filing by small entity, if applicable. A Small Entity Statement must also be filed (Note 37 C.F.R. 1.9, 1.27, 1.28).

SUBTOTAL =

Processing fee of \$130.00, for furnishing the English Translation later than ☐ 20 ☐ 30 months from the earliest claimed priority date (37 C.F.R. 1.492(f)).

TOTAL NATIONAL FEE =

Fee for recording the enclosed assignment (37 C.F.R. 1.21(h)). The assignment must be accompanied by an appropriate cover sheet (37 C.F.R. 3.28, 3.31). \$40.00 per property +

Fee for Petition to Revive Unintentionally Abandoned Application (\$1240.00 - Small Entity = \$620.00)

TOTAL FEES ENCLOSED =

CALCULATIONS PTO USE ONLY

\$ 860.00

\$ 0.00

\$ 0.00

\$ 160.00

\$ 0.00

\$ 1020.00

\$ 0.00

\$ 1020.00

\$ 0.00

\$ 1020.00

\$ 40.00

\$ 0.00

\$ 1060.00

Amount to be:

refunded \$

Charged \$

- a. ☒ A check in the amount of \$1060.00 to cover the above fees is enclosed.
- b. ☐ Please charge my Deposit Account No. 14-1140 in the amount of \$_____ to cover the above fees. A duplicate copy of this form is enclosed.
- c. ☒ The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment to Deposit Account No. 14-1140. A duplicate copy of this form is enclosed.
- d. ☐ The entire content of the foreign application(s), referred to in this application is/are hereby incorporated by reference in this application.

NOTE: Where an appropriate time limit under 37 C.F.R. 1.494 or 1.495 has not been met, a petition to revive (37 C.F.R. 1.137(a) or (b)) must be filed and granted to restore the application to pending status.

SEND ALL CORRESPONDENCE TO:

NIXON & VANDERHYE P.C.
1100 North Glebe Road, 8th Floor
Arlington, Virginia 22201
Telephone: (703) 816-4000

SIGNATURE

Larry S. Nixon

NAME

25,640

REGISTRATION NUMBER

November 6, 2000

Date

09/674720

529 Rec'd PCT/PTC 06 NOV 2000

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of

BRISCOE et al.

Atty. Ref.: **36-1385**

Serial No. **(To Be Assigned)**

Group:

National Phase of **PCT/GB99/01773**

Filed: **November 6, 2000**

Examiner:

For: **COMMUNICATIONS NETWORK**

November 6, 2000

Assistant Commissioner for Patents

Washington, DC 20231

Sir:

PRELIMINARY AMENDMENT

Prior to calculation of the filing fee and in order to place the above identified application in better condition for examination, please amend the claims as follows:

IN THE CLAIMS

Claim 6, line 1, delete "or 5",

Claim 7, line 1, change "any one of the preceding claims" to --claim 1--,

Claim 8, line 1, change "any one of the preceding claims" to --claim 1--,

Claim 9, line 1, change "any one of the preceding claims" to --claim 1--,

Claim 11, line 1, change "any one of the preceding claims" to --claim 1--,

Claim 12, line 1, change "any one of the preceding claims" to --claim 1--.

Change claim 14 to claim 16.

Change claim 15 to claim 17.

REMARKS

The above amendments are made to place the claims in a more traditional format.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By:



Larry S. Nixon

Reg. No. 25,640

LSN:ms

1100 North Glebe Road, 8th Floor

Arlington, VA 22201-4714

Telephone: (703) 816-4000

Facsimile: (703) 816-4100

17/PRTS

529 Rec'd PCT/PTC

09/674720

06 NOV 2000

Communications Network

BACKGROUND TO THE INVENTION

The present invention relates to a communications network, and in particular
5 to charging mechanisms in such a network. It includes aspects of the inventions
disclosed and claimed in the present applicant's co-pending British patent application
no. 9812161.9 filed 5 June 1998 (applicant's ref: A25547) and the contents of that
earlier application are incorporated herein by reference.

In conventional communications networks, such as national PSTNs (public
10 switched telephone networks), a significant proportion of the network resources are
devoted to metering and billing network usage. Studies have estimated these
resources as consuming as much as 6% of the revenue of a telecommunications
company. The Internet, by contrast, does not in general incorporate metering and
15 billing mechanisms for individual customers. The absence of the network
infrastructure required to support metering and billing reduces the operational costs
of the Internet compared to conventional telephony networks, and has facilitated the
rapid expansion of the Internet. However, the absence of appropriate billing
mechanisms has significant disadvantages in terms of the characteristics of the
20 traffic carried by the Internet: it encourages profligate use of network resources, and
diminishes the incentive for investment in network infrastructure to support new
applications requiring, e.g., guaranteed quality of service (QoS).

SUMMARY OF THE INVENTION

According to a first aspect of the present invention, there is provided a
25 method of operating a communications network, including automatically varying,
depending on network loading as detected at a customer terminal, a tariff for
network usage by the customer terminal.

In this document references to varying a tariff encompass changing a price,
e.g. as a function of a congestion parameter. The structure of the tariff may be
30 unchanged.

The present invention provides an approach to charging for network usage
which involves little or no network overheads, and which moreover is able to reflect
local variations in network loading. This is achieved by detecting a measure of

network loading at the customer terminal. This measure will reflect those resources being used by the customer terminal at a particular time. For example, in a federated data network comprising a number of sub-domains, such as the Internet, a customer terminal in one subdomain may attempt to access a server in another subdomain at a time when there is overloading of a router somewhere in the path from the customer terminal to the server. Then, even if average loading of the network as a whole is low, the loading as perceived by the customer terminal is high, and a tariff for use of the network by the customer terminal is increased to reflect the scarcity of the relevant network resources. Similarly other terminals attempting to access data via the same router will also perceive the network loading as being high, and will also have their tariffs increased accordingly. With an appropriately steep rise in tariff, at least some customer terminals will then choose to abandon or defer their attempts to use the network resource in question, thereby relieving the loading of the router.

This aspect of the invention may be used in systems in which the end user is merely informed of the price and accounting and billing is carried out by the network provider, also in systems where the end user measures their usage and supplies this information to the network provider, and also in systems where the end user both measures their usage and calculates the payment due to the network provider.

The method may include detecting at the customer terminal a network performance parameter which depends on network loading, and varying the tariff depending on the network performance parameter. When the network is a packet network, then the network performance parameter may be the number of packets lost in transmission between a data source and the customer terminal. This approach has the advantage that it is suitable for implementation using existing congestion control mechanisms in packet networks. For example, in the Internet, routers respond to congestion by dropping packets, and a customer terminal using TCP (Transport Control Protocol) detects packet loss and controls a congestion window at the terminal accordingly. The present invention may use the detected packet loss to trigger an increase in the tariff applicable to the terminal. This may be accompanied by the signalling of congestion by the receiving terminal to the data source.

In an alternative and preferred approach, when congestion is detected within the network an explicit congestion signal may be generated and transmitted to the customer terminal, and the receipt of the explicit congestion signal at the customer

terminal may then trigger the increase in the tariff. This approach has the advantage that an appropriate congestion signal may be generated in advance of significant degradation in network performance occurring, making it easier to maintain a required quality of service. Preferably the explicit congestion signal is carried with a data packet on the communications network. Preferably a router in the network writes an explicit congestion signal in a packet when congestion is detected at the router.

As a further alternative, the congestion notification signal may be a source quench signal generated at the router and transmitted back to the data source.

Preferably there is a non-linear relationship between the increase in tariff and the detected network loading. Preferably the method includes making a first relatively smaller increase in the tariff when congestion is first detected, and making at least one further, relatively larger increase, if the congestion persists. This behaviour may be pre-programmed into the tariff algorithm. In this way oscillatory behaviour of the network can be avoided, while ensuring that a sufficient increase is made to have a desired impact on the demand for network resources.

Preferably the method includes programming a decision agent at the customer terminal with user-determined price criteria, and comparing a price calculated using the tariff with the said price criteria. For example, the decision agent might be programmed to proceed with data communications as long as a per-packet charge is zero or is less than a predetermined threshold. Then if congestion causes the charge to rise above that threshold, the data communication is interrupted and the decision agent modifies the operation of the customer terminal. For example the decision agent may slow down communications once the cost threshold is exceeded. The user may then choose to continue with the communication, or may elect to abandon the attempt to access data from a particular source. The decision agent may be programmed with rules which relate to an overall price for multiple simultaneous transmissions and which accord different weights to different respective applications.

Preferably the method includes distributing a tariff algorithm via the communications network to a plurality of terminals and calculating at each terminal using the tariff a charge for network usage by the terminal. In this case preferably the method includes steps, carried out by a network operator, of intermittently

sampling traffic between a customer terminal and the network, and as part of the sampling recording network loading affecting the customer terminal;

for the sampled traffic comparing a charge calculated by the customer terminal and an expected charge and detecting thereby any discrepancy.

5 The step of recording network loading may involve, for example, recording packet drops, or explicit congestion notification (ECN) bits or source quench messages.

The invention also encompasses a network operating by a method in accordance with the first aspect of the invention, and customer terminals, network
10 routers and other network nodes for use in such a network.

DESCRIPTION OF DRAWINGS

Systems embodying the present invention will now be described in further detail, by way of example only, with reference to the accompanying drawings, in
15 which:

Figure 1 is a schematic showing a network embodying the invention;

Figures 2a and 2b are graphs showing tariff functions;

Figure 3 shows the format of a differential service byte;

Figure 4 is a schematic showing the component objects of a charging
20 architecture for use with the network of Figure 1;

Figure 5 shows data passed between the accounting objects of Figure 4;

Figure 6 is a schematic showing protocol stacks on a customer terminal and in the network domain;

Figure 7 is a graph showing the variation of tariff with time;

25 Figures 8a to 8e are class diagrams for software implementing accounting and measurement objects;

Figure 9 is a diagram showing a graphic user interface (GUI) for use with the objects of figures 8a to 8e;

Figure 10 is a class diagram for software implementing tariff objects; and

30 Figure 11 is a diagram showing an alternative embodiment.

DESCRIPTION OF EXAMPLES

As shown in Figure 1, a communications network 1 includes a number of network sub-domains 2A-C. The network sub-domains may be under the control of different operators who may not trust each other. The network subdomains are interconnected by gateway routers 3, 4. In the present example the communications network is the Internet and supports both unicast and multicast Internet Protocol (IP) and associated protocols. A customer terminal 5 is connected via a public switched telephony network (PSTN) 6 and an access router 7 to a subdomain 2A. A single blocking test is applied to traffic at this point of access. The gateway routers 3,4, and access router 7 may be commercially available devices such as CISCO series 7500 routers and CISCO series AS5800 universal access server respectively. Other customer terminals are connected to the network, including a Java-enabled mobile terminal 8 and a data server 9. The customer terminal 5 may be connected via a LAN to an accounting server. The accounting server may include an accounting object as described below that receives measurement data from the customer terminal.

In addition to the local tariff variation mechanism that is described below, the network also uses network-based control of a number of tariff bands. A network management platform 10 is connected to each subdomain. Each network management platform may comprise, for example, a computing system comprising a SPARC workstation running UNIX (Solaris) together with network management applications. The network management platform 10 hosts management entities and tariff entities. The network management platform communicates with agents 100 in managed devices connected to the respective subdomain, for example using SNMP (simple network management protocol). The management platforms monitors the overall loading of network resources in the respective subdomains, and, as will be further described below, adjust the tariffs for network use accordingly. The Net management platform (NMP) instructs the agent to monitor the device and report aggregated results at regular intervals back to the NMP, so the NMP can monitor the combination of all reports.

Tariff data is communicated to peer tariff entities in other subdomains and also to the customer terminals. The tariff data is multicast using, for example Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) dense mode. The tariff data channels are announced and monitored

using protocols based on SDP (Session Description Protocol), SAP (Session Announcement Protocol) Charging is carried out on a "pay and display" model. Each customer terminal monitors its own network usage, for example by counting the number of packets it sends or receives across the network interface and the quantity of data (in bytes) in those packets. It calculates, using a tariff received via the network, the payment due to the network operator, and makes a corresponding payment into an account at the network operator. The network operator polices the use made by customers of the terminal by intermittently sampling traffic to or from a particular customer and comparing the use made and the use paid for.

The tariffs supplied to the customer terminals are divided into bands of different volatilities. The tariffs are varied under the control of the network operators to reflect the overall loading of the network. That is to say, if network loading becomes high, then the tariffs may be increased to reflect the scarcity of network resources.

A service provider may offer different products defined by different service level agreements, and/or by different price volatilities. For example product A might offer best-effort service at a fixed price while another product B might offer best-effort service at a variable price. A service provider may adjust product prices on the basis of the following parameters: the price the service provider pays to its wholesale provider: competitors' prices; current resource utilisation; relevant demand for different products. In response to changes in these parameters, tariff adjustments may be effected in one of three ways. Firstly, a tariff may adjust prices on the basis of local observations of network loading, without necessitating explicit communication from the provider. This approach, which is described in further detail below, needs to be built into the tariff at the outset, and is limited to those price variations which are dependent exclusively on local observations. Secondly, the provider may tune a tariff by adjusting some of its parameters. This kind of adjustment is required when the decision is dependent on parameters which cannot be observed directly by the customer, e.g., variation in the wholesale price of network resources. Thirdly, the provider may completely replace a tariff. This is required when the existing tariff cannot accommodate the changes that are required.

The first of the tariff changes described above is necessarily carried out automatically. The second type of change may be performed manually, or by an

00301-0242960

agent that issues adjustments automatically in response to observations made by the service provider system. The third type of change is likely to be performed manually, as replacement of a new tariff will in general require an element of design requiring human input. However, it is possible that an agent might be employed to

5 automatically switch tariffs for a product on the basis of a set of specified rules.

This section describes a prototype that we implemented to demonstrate the tariff subsystem outlined above. Features of the design include:

- using mobile code to represent tariffs and associated user interface components;
- use of a repeated multicast announcement protocol to communicate tariffs and
- 10 tariff adjustments efficiently;
- using dynamic class loading and reflection in order to receive and tune tariffs.

The prototype consists of a library of general-purpose Java classes and two specific applications, namely:

- a provider system which allows the provider to introduce, replace, and tune tariffs
- 15 for a number of products;
- a customer system that enables customer to keep track of the charges being applied for the products they are using.

The provider system services multiple instances of the customer system running on different hosts in a multicast-enabled network. A multicast announcement protocol is

20 used to communicate tariff changes from the provider system to customer systems.

In order to maximise flexibility with respect to the definition of tariffs, we chose to represent tariffs using Java classes. This technique is also used to supply user interface components to customers to support visualisation of the behaviour of a tariff.

25 The `Tariff` interface acts as the base class for all tariffs. This defines a single operation `get GUI()` which returns as a Java SWING component that can be incorporated into the customer's GUI (graphical user interface). This GUI component enables the customer to visualise the behaviour of the tariff using techniques appropriate to the tariff.

30 Subclasses of the `Tariff` interface establish a set of tariff types, each of which is associated with a different set of measurement and input parameters. These parameters are identified by listing them in the signature of the `getCharge()` method. For example, the interface `RSVPTariff` defines `getCharge()` as

receiving n RSVP TSPEC, allowing for the definition of tariffs that compute price on the basis of the characteristics of an RSVP reservation. On the other hand, the interface PacketCountTariff defines getCharge () as receiving measurements of packets in, packets out, and current congestion (typically measured as a function of
5 packet drop), allowing for the definition of tariffs that are dependent on packet counts and sensitive to congestion. Other tariffs may be added as new forms of usage-measurement emerge.

Tariffs are defined by providing implementations of the various tariff interfaces described above. For example, the tariff PacketCountLinear implements
10 PacketCountTariff to compute charges in proportion to packet counts. Another tariff CongestionSensitiveLinear works on a similar basis, but adds a penalty charge if the customer does not stay within a specified traffic limit in the presence of congestion.

In addition to the tariff interface implementation, a tariff may make use of
15 other 'helper' classes to assist in its operation, as well as one or more user interface component classes for customer visualisation purposes. A provider-side user interface may also be required in order to enable the provider to make tariff adjustments.

A complete tariff description consists of a set of Java classes, some of
20 which are destined for the customer system and others which are intended for use by the provider system. The customer-side classes are bundled into a Java archive (JAR) file to facilitate processing by the provider system.

In order to deploy a new tariff, the provider system first loads the tariff classes which it requires into its execution environment. It then loads the customer-
25 side bundle, serialises it, signs it with a private key, and uses an announcement protocol to distribute it to customer systems. The use of a signature makes it possible for customers to verify that received tariffs are authentic.

Upon receiving the bundle, each customer system verifies the signature (using the public key matching the provider's private key), and at the activation time
30 specified in the announcement protocol headers which may be significantly later, e.g. hours or days, unpacks the bundle, and loads the classes into its execution environment using a purpose-built dynamic class loader. An instance of the received tariff class is created and installed in place of the previous tariff. If the tariff has a

user interface component (obtained by calling the tariff object's `getGUI ()` method), then it replaces the user interface of the previous tariff. The change in user interface serves to notify the user that the tariff has changed.

Tariff adjustment involves the remote invocation of an operation which is
5 specific to the tariff currently in force. This means that a customer system cannot know the signature of this operation in advance of receiving the tariff i.e. the operation will not be listed in any of the tariff interfaces known to the customer system.

In order to get around this problem, use is made of the "reflection" feature
10 supported by Java. In order to disseminate a tariff adjustment, the provider creates an instance of an `Invocation` object, which stores the name of the operation to be called, together with the parameters that are to be supplied to it. This object is then serialised, signed, and announced using the announcement protocol. When an adjustment is receive and verified by a customer system, the `Invocation` object is
15 de-serialised and applied to the current tariff by using reflection to invoke the described operation.

In order to simplify the announcement protocol, adjustments are required to be idempotent and complete. Idempotency guarantees that a tariff will not be adversely affected if an adjustment is applied more than once. Completeness implies
20 that an adjustment determines the entire parameter set of a tariff object, so that an adjustment completely removed the effects of any previously applied adjustments.

The customer system may apply a tariff by repeatedly invoking the `getCharge()` operation supported by that tariff every second, and adding the returned value to the cumulative charge. The parameters supplied to `getCharge`
25 `()` depend on the kind of tariff currently in force. For example, if the tariff is an implementation of `PacketCountTariff`, then measurements of inbound packets, outbound packets and congestion over the past second are required. However, if the tariff is an implementation of `RsvpTariff`, then only a `TSPEC` describing the current reservation is required. This implies that a customer system can only subscribe to a
30 product if it can supply the parameters require by the tariff associated with hat product.

Each invocation of the `getCharge ()` method also results in an update to the tariff-specific user interface. For example, in the `CongestionSensitiveLinear` tariff, the usage parameters supplied to `getCharge ()` are used to update the graphical displays of traffic and congestion.

- 5 The announcement protocol is used to communicate serialised tariffs and adjustments from a provider system to multiple customer systems. The number of customer systems is assumed to be large, and a repeated multicast solution is adopted.

- 10 Each product supported by a provider is assigned a multicast channel for announcement purposes. Customer systems listen to the channels corresponding to the products that they are using. In the current implementation, it is assumed that each customer system has knowledge of well-known multicast addresses for the products it is interested in.

- 15 For each product channel, the provider repeatedly announces the current tariff and the most recent adjustment made to it (if any). Each announcement carries a version number, which is incremented each time the announcement is changed. Customer systems only process announcements when a version number change is detected. If a new customer joins a channel, it waits until it receives a tariff before processing any adjustment announcements. Furthermore, an adjustment is only
20 applied if its announcement version is greater than that of the current tariff, thereby ensuring that a missed tariff announcement does not result in the application of a subsequent adjustment to an old tariff.

- 25 While centralised monitoring and control of tariffs by the network management platform is effective to respond to global changes in the loading of the network, it is difficult to handle localised congestion in this way. It is difficult to cause a price rise signal to be multicast in such a way that the signal is only received by those attempting to communicate packets through the point of congestion. This would require a separate multicast transmission for each element in the Internet , e.g. a multicast for every different queue on every interface of every router. Alternatively
30 some aggregation of price rises triggered by local resource loading might be used. This however would mean that price rise signals were sent to users who were not making use of the congested resource. This in turn would make it necessary for the

To overcome these difficulties, the tariff algorithms installed on the customer terminals are arranged to respond automatically to congestion in a network resource being used by the terminal. Each algorithm includes a function which varies the price for network usage in dependence upon the detected congestion level. This function may be integrated in the main tariff algorithm, or, as in the example described here may be a separate algorithm used to calculate a premium to be added to a price calculated in accordance with the main tariff algorithm.

The main tariff algorithm calculates a price P as a function of a number of quality parameters, Q_1 , Q_2 , Q_3 where, for example, Q_1 is a specified latency for packets communicated across the interface between the customer terminal and the network, Q_2 is the reserved bandwidth for the transmission, Q_3 is a specified level of reliability corresponding to a maximum permissible level of packet loss.

The price P is then given by:

$$P = f(Q_1, Q_2, Q_3, \dots)$$

An example of the pricing function in terms of one of the quality parameters Q is shown schematically in Figure 2a.

The congestion tariff algorithm calculates a premium ΔP which is a function of one or more congestion parameters C:

$$\Delta P = f(C_1, C_2 \dots)$$

The congestion parameters provide a measure of the loading of the resources which a customer terminal is making use of at any given time. In the present example the ratio of packets lost to packets received is used as the congestion parameter. This parameter is readily calculated, for example in the case of packets using TCP (transport control protocol), or RTP (real time protocol) over UDP (user datagram protocol), since such packets include a sequence number. Figure 2b shows one example of the function for generating the premium. In this case, the premium increases as an approximately exponential function of the congestion, so that at low congestion levels a small premium is charged, while if congestion increases still further, then at higher levels of congestion the premium increases sharply.

In an alternative implementation, an explicit congestion signal is added by any congested router within the network to packets transmitted to the customer

terminal. Whether congestion is signalled explicitly or implicitly by packet drop, the price algorithm might not directly charge for congestion indication. Instead, it may charge if the customer terminal does not halve the congestion window in response to a congestion indication. Halving the congestion window halves the rate of transmission which is the standard response to congestion indication used in TCP and other compatible congestion control protocols. The change in transmission rate can be measured by the network provider as well as at the customer terminal, although this places a heavy load on router processing resources. If the provider is merely sampling certain customers to audit their own measurements, this results in an acceptable load.

Although only a single main tariff and premium are described here, in practice different subdomains, and different service providers associated with each subdomain, may each have a different pricing structure, with different main and premium tariffs. However, there is across all the subdomains a common relationship between network loading levels and congestion signalling.

The operation of this second implementation will now be described in the context of a network operating using a differentiated service as described in the Internet Engineering Task Force draft "Differentiated Services Operational Model and Definitions" and in the paper by David D Clark (MIT), "A Model for Cost Allocation and Pricing in the Internet", presented at MIT Workshop on Internet Economics, Mar 1995, URL:<http://www.press.umich.edu/jep/works/ClarkModel.html>. In a network implementing differentiated services, nodes are arranged to discriminate between packets to provide different levels of service. This capability might be used, for example, to accord delay-sensitive data, such as data generated by an IP telephony client, a higher priority compared to other data, such as email data. At the network edge, for example at a client terminal running the IP telephony client, bits in a TOS (type of service) octet contained within each packet header are set to indicate the appropriate service level. Those bits are used by routers within the network to determine how the relevant packets should be handled.

The TOS octet when used in this way is termed the DS (differential service) byte. The format of the differential service byte is shown in Figure 3. Bit zero, labelled "IN" indicates whether the packet is inside or outside a defined profile. Bits 1 to 5 labelled "PHB" define a "per-hop-behaviour" that is they indicate how, for

example, a router should handle the packet, e.g. by according it lower or higher priority. Bits 6 to 7, in this particular form of the DS byte, are used for explicit congestion notification (ECN). One of these bits is set to indicate whether the routers in the path of the packet are capable of setting the ECN field, and the other bit is used as a flag which is set (by ECN capable routers) when congestion, or loading which would potentially lead to congestion, occurs. Random Early Detection (RED) algorithms are currently implemented in routers. These algorithms measure average queue length within the packet buffers of a router. An exponential moving average is calculated. When that average queue length exceeds a predetermined threshold, then the router signals that congestion is occurring. Conventionally this signalling has been done simply by dropping a packet. However, in the context of an ECN scheme, the router, instead of dropping a packet, sets an ECN bit in a packet header to indicate that congestion is occurring. This is done probabilistically: that is, some only of the packets passing through the router are marked. The probability of a packet being marked increases with the average queue size. In the rare case that the queue increases to a length where the router buffers are full, then packets are dropped, rather than an ECN bit being set. In this case ECN bits are set for all the remaining packets.

In operation, if the client terminal 5 is accessing a data source on the server 9, congestion may occur, for example, at router 4 which links network sub-domains 2B and 2C. RED-like algorithms in the router 4 detect that the queue lengths in the router buffers, as calculated using the exponential moving average, exceed a predetermined threshold. Accordingly some of the packets from the server 9 en route to the client terminal have the ECN bit of the DS byte set by the router 9 to mark the fact that congestion is occurring. At the client terminal, the DS byte in the headers of incoming packets is read. A moving average of the number of packets containing an ECN bit which is marked is calculated. This average then provides the congestion parameter C_1 which is used to calculate the premium:

$$\Delta P = f(C_1).$$

The total price to the user P_{TOT} is then calculated by adding together the prices determined by main tariff algorithm and by the premium algorithm:

$$P_{TOT} = P + \Delta P.$$

This total price is passed to a cost decision agent running on the client terminal. This cost decision agent is programmed with user defined rules. These might state, for example, that the cost decision agent should authorise the system to proceed with a connection as long as the total cost averaged over a certain time period falls below a
5 predetermined threshold, e.g. of £0.01 per minute, and that the cost decision agent should suspend a connection and alert the user if the cost rises above that threshold. Alternatively, as previously noted, the cost decision agent may handle several applications simultaneously, and may be programmed to slow down one of the applications as the premium for using a data source accessed by that application
10 increases.

For some customers in some circumstances, the variability of the tariff for network usage may pose problems. For example, a customer might wish to schedule a multicast audio-video conference with, say, twenty participants and with the customer intending to pay the network usage costs for all of the participants. The
15 conference will be scheduled a day or more in advance, but since the tariff is variable over time, at the point when the conference is scheduled, the customer will not know the cost which will be incurred. The preferred implementation of the present invention overcomes this problem for the customer, and also provides an additional revenue source for the network operator. This is achieved by applying a further tariff
20 which defines a premium for price stability for a given customer terminal, the value of the premium increasing as the required duration of the period of stability increases. In the example referred to above, the customer scheduling a conference one day in advance refers to the current tariff to determine the present cost of the network resources required, including any congestion premium currently applicable, and then
25 refers to the further tariff to determine the premium for maintaining the cost at its present value for a period of one day. Then if the user decides to make use of the fixed tariff a message is sent from the user to the network operator. The message identifies the data flows in question, for example by specifying the IP addresses and the service levels for the audio and video data flows between the participants,
30 together with the current time at which the user is contracting for a fixed period, and the required duration of the fixed period and also an identifying code for the message. Alternatively, the message with these details may be tracked with an appropriate cryptographic function (e.g. MD5), and the hash value only communicated initially to

the provider, and subsequently used to verify account data. Subsequently, for the duration of the specified period, in this case one day, the user applies the fixed tariff. In systems where the users themselves are responsible for generating accounting data and calculating the applicable charge for network usage, then the user calculates
5 the charge using the fixed tariff and returns the accounting data to the network operator together with the identifying code for the fixed tariff contract.

This aspect of the invention is not only applicable to systems where tariff variation is carried out locally in response to the detection of congestion, but may also be used in systems where, for example, the only variation is imposed by the
10 network operator.

The tariff for the stability premium may be distributed by multicasting from tariff entities in network operator platforms to the customer terminals in the same manner as described previously for the other tariffs discussed above. As with those other tariffs, different tariffs and optionally different periods of stability might apply
15 to different service levels. In a scenario such as that discussed in the preceding paragraph, then different service levels might be used for different elements of the audio and video data streams required for the multicast conference.

Figure 2c shows a possible for the stability premium tariff. This tariff may be embodied in a Java function multicast to the users. Curve C is communicated to the customer terminal as another subsidiary algorithm that contributes to the main tariff
20 algorithm. The customer may choose a period of price stability and calculate the premium above the spot price that his will require using curve C. Curve B is an example of the stable price chosen in a particular case. Once the period of stability expires a new one can be bought at the premium over the spot price at that time. To
25 buy a period of price stability, the customer must announce the period required and the range of addresses for which it applies to her provider. It may be required for one address (e.g. for the duration of an Internet phone call to a single person) or for a range of addresses (e.g. for a video conference). Certain risk-averse customers might request stable pricing for all addresses all the time.

30 For ease of description, the preceding sections have treated in isolation the local variations in tariff in response to congestion. In practice, this mechanism will in general be combined with other responses to congestion, and with other sources of variation in the tariff. Also, a decision to proceed with a transmission despite

00674730 449909

- congestion will in general require the consent of parties at both ends of the transmission. Considering the entire system of the data source, network and routers and the data receiver, the implementation of an increase in tariff (also termed here a "fine") in response to locally detected congestion occurs as a last resort. Other
- 5 responses are implemented first, in the following numerical order:
- 1.the network re-routes around congestion
 - 2.the network borrows capacity from lower levels of service (lower in the context of the relevant dimension(s) of QoS) including the best effort service
 - 3.the network introduces extra capacity (possibly automatically)
 - 10 4.the end-system establishes that the congestion is on the shared network and not just on the access links or end systems
 - 5.the end-system sets QoS requirements to a "higher" level (if cheaper than the fine for ignoring congestion at the current level)
 - 6.the end-system decides it is essential to ignore the congestion, given the fine for
 - 15 doing
 - so might be quite high
 - 7.both (all) end-systems agree to ignore the congestion.

Typically, it is at step 4 that an ECN signal is generated. Steps 1 to 3 precede the generation of this signal and steps 5 to 7 follow the generation of the

20 ECN signal.

The last step prior to proceeding with a connection and paying the premium for doing so is establishing agreement by both parties. Accordingly, when the customer terminal detects congestion, either through receiving explicit congestion notification, or through detection of a relevant parameter such as packet loss, the

25 customer terminal signals this fact back to the or each other end system. In the present example therefore, the client terminal 5 signals to the data server 9 that congestion is occurring. The data server is programmed with rules, which as at the customer may be implemented as an agent, which determine the response to such a signal. For example, the server may refuse service in these conditions. As

30 described previously with respect to Figure 1, in the present example tariffs are multicast through the network from network operators to the customer terminals, and charging is carried out using a "pay and display" process. Figures 4a and 4b shows the objects used to implement the charging architecture in this case. Figure 4a

shows the higher level objects and 4b shows the component objects used in a software implementation of the architecture of Figure 4b. In Figure 4a, objects on the client terminal are shown in the half of the Figure labelled "customer" and objects on the access router 7 and the corresponding network sub-domain are shown in the half of the Figure labelled "edge network". The objects on the customer terminal include a session control object S , a customer business rules object B_c , a customer pricing object Pr_c , a QoS manager Q , a customer accounting object Act_c and a customer measurement object M_c . The business rules object B_c receives information on those aspects of the session which involve liability for payment and receives current pricing data from the pricing object Pr_c . The customer business object makes decisions, under the customer's policy control on which chargeable services are utilised, and how much of the chargeable services are utilised. These decisions are fed to the QoS manager Q , which decides which mechanisms are used to achieve the requirements. The QoS manager then controls the customer measurement object M_c to determine which aspects of traffic and service to measure and which aspects to ignore. The measurement object then records the selected aspects of the traffic, for example counting the number of packets received by the customer terminal and the QoS levels for those packets. These data together with the current tariffs, including any premium for congestion, are then used by the customer terminal to determine the charge payable to the network operator. The measurement object M_c is also programmed with instructions which determine the frequency with which it passes data to the customer accounting object Act_c . The customer accounting object Act_c passes payments to an accounting object Act_p in the network provider's domain.

The accounting objects on the customer terminal may be implemented using a small encrypted flat-file database. On the network provider's side, the equivalent objects may be implemented using a larger database that is scaleable to handle e.g., tens of thousands of customer accounts. An object request broker (ORB) is used for communication between the customer-side objects and the network-side objects, implemented using commercially available tools such as ORBIX (TradeMark) from Iona Technologies plc.

On the network provider's side, that is to say within the subdomain to which the customer terminal is connected, the customer's traffic is measured by a version of M , denoted M_p , but only on a sampling basis determined by the policing function,

Po. That is to say, the network operator samples the customer's traffic only intermittently. Po controls where in the network measurements are made in order to capture all of any particular customer's traffic. A bulk measurement function, M_b , is responsible for reporting aggregate traffic levels, as reflected in the moving average of the router queue lengths, to the pricing object, Pr_p . Bulk measurements would typically be collected from across the provider's domain to a centralised pricing function (which would be replicated for reliability). Pr_p sets prices taking into account the business rules from the network provider's business object, B_p , as well as the current traffic levels reported by M_b and pricing from neighbouring providers (see below). The policing function, Po , compares sample measurements from M_p with accounting messages received at Act_p as a result of the customers own measurements. If it establishes that the accounts are insufficient it might restrict service at the access control gateway, Acs , or initiate some other punishment. Encapsulated within the accounting object another policing object checks the accounts match the payments within the contracted time for payment. Finally, the identity mapping function, I , provides a mapping between a customer's identity (account, digital signature, etc.) and their current network address (typically allocated by the ISP, whether unicast or multicast).

Figure 5 shows the data which are passed between the accounting objects. In this example the account data comprises: account identity; bill record identity; service type identifier; source address; destination address; tariff identity; time; period (i.e. the period covered by the bill record); units; cost; and currency. In addition, the payment data comprises the amount of money and the currency of payment.

Figure 6 shows the measurement region within protocol stacks on the customer terminal and in the network domain. Ideally there would be two measurement points within this region, one trusted by the customer and one trusted by the network, for example at the two points referenced (a) in the Figure. For ease of implementation, a single measurement point (b) trusted by both parties may be used. This might be implemented, for example within a secure module such as a cryptographic card on the client terminal. As an alternative, measurements may be made at different points with some possibility of discrepancies between measurements. On the network the practical measurement point is at the first

access device(s) that, for each customer, inspects network layer headers (c)(IP in this case). ISPs should not measure any deeper into their network (d) because their access network and systems will introduce delays and losses.

For an individual customer (e.g. on dial-up access), a practical point at which to measure would also be alongside the network layer but in their end-system's stack (e). Ideally these measurement points would be lower in each stack to be closer to the interface between the two parties and less likely to be affected by contention in the stack. However, measuring at the link layer (f-f) would be inappropriate because only some chargeable parameters set at the network layer will ever be reflected in link layer frames; network level multicast, end-end latency requirements etc. may never be visible at the link layer. Also, link layer headers would need to be ignored when measuring packet sizes for bandwidth calculations to avoid apparent discrepancies where different link technologies are chained together.

In the reception direction (up the stack) this choice of measurement points implies that the lower layers must be dimensioned (buffer sizes, interrupt and thread scheduling priorities) to cope with the most stringent QoS requirements of higher layers. As frames are taken off the physical media, the machine must be able to pass data up the stack without any chance that usage-charged data gets discarded (e.g. due to buffer overflow caused by interrupt contention) before it gets to the network layer. It is at the network layer where the ISP's service is to be measured and where it is most convenient for QoS requirements to control correct differential treatment of the various flows as they are passed further up the stack (on end-systems) or forwarded (on routers).

The measurement objects described above may be implemented using, with appropriate modifications, publicly available network metering software such as Nevil Brownlee's NeTraMet system. This is a software meter which conforms to the IETF internet accounting architecture described in RFC 2063 and RFC 2064. The meter builds up, using "packet sniffing", packet and byte counts for traffic flows, which are defined by their end-point addresses. Although generally, Addresses can be ethernet addresses, protocol addresses (IP, DECnet, EtherTalk, IPX or CLNS) or 'transport' addresses (IP port numbers, etc), or any combination of these, in the present implementation IP addresses only are used. The traffic flows to be observed are specified by a set of rules, which are downloaded to NeTraMet by a 'manager'

program. Traffic flow data is collected via SNMP (Simple Network Management Protocol) from a 'collector' program.

Figure 7 shows how the main tariff determined by the network operator varies in time. In the Figure, curve A is the spot price calculated to reflect the loading of the network at any instant. Curve B is one of a number of different tariff bands. Different tariff bands have different volatilities, and the customer pays a premium for bands offering greater stability. Tariffs are communicated to the customer terminals using a hierarchy of channels carried by the network. An initial contract between a customer and a service provider a single channel address that might typically hold new announcements distributed some months apart (e.g. for contract variations or for new services specifying which second level channel to listen to for tariffs or for downloading new code to handle new tariff structures). The second level channels might deliver updates hours apart which simply announce the addresses of third level channels for the most volatile information. These third level channels may carry updates at intervals of less than a second. Prices for many services may be carried on one channel. For greatest efficiency, this one channel may be split into several channels at times of highest volatility, and re-aggregated into a single channel in more stable periods.

Tables 1 to 7 below list Java source code used to implement two different tariffs. The code of table 1 establishes the operations used for communication between a customer system and a tariff algorithm downloaded by the customer system. Table 2 shows a linear tariff algorithm, in which the tariff depends on the total of the packets sent and packets received by the customer together with a congestion parameter. Table 3 shows the code for generating the customer display in this case. Table 4 shows the code used to display the tariff at the network operator's server. Table 5 shows an exponential tariff algorithm. Table 6 generates the customer display and Table 7 the operator display for the exponential tariff algorithm. By downloading Java code to generate the user interface, that interface can be tailored to the requirements of the particular tariff, and can be adapted as the tariff changes.

Although the examples so far described have been in the context of federated packet data networks, such as the Internet, many aspects of the invention can also be used with advantage in other types of network, such as in a circuit-

switched PSTN (public switched telephony network). Figure 11 shows an example of the invention applied in this context. In this network, customer terminals 81, which are in this example so-called intelligent phones, that is telephones incorporating a microprocessor and a data interface, are connected via local exchanges 82 and trunk exchanges 3 to the telephony networks. The trunk exchanges 83 are connected via a common channel SS7 (signalling system number 7) signalling network to a service control point 85 that is responsible for the execution of advanced call control functions. The service control point 85 is also connected to an operational support server 86 that is responsible for billing operations, and that, in this example, controls the setting of tariffs for the network. The OSS server and customer terminals include tariff entities (TE). The fixed PSTN network is also interconnected via a gateway 87 to a cellular GSM network 88. Base Stations BS in the cellular network communicate signals to intelligent mobile phones 89. In operation, network tariffs are distributed to customer terminals via the PSTN network and via the GSM network. Conveniently, the tariff may again take the form of Java functions which are executed on processors in the customer terminals. The Java functions may be streamed as Internet packets. In one implementation, these Internet packets may be distributed via the PSTN networks and GSM networks themselves. For example, the packets may be encapsulated and transported to the trunk exchanges using the MTP (message transport part) transport layer and may be communicated onwards to the customer terminals using out-of-band signalling. Alternatively, a separate data connection may be established between the OSS server and the customer terminals via the public internet. As in the examples above, the network operator monitors the loading of resources within the network and may transmit signals to the tariff entities in the customer terminals to change the tariff to reflect the scarceness or otherwise of relevant resources. Customer terminals may themselves monitor network loading and automatically generate variations in the tariffs. Usage of network resources may be measured locally by the customer terminals instead of conventional billing carried out within the network. The network operator may police the measurement of usage data by carrying out sampling, as described previously.

Table 1

5	<pre>// Generated by Together package com.bt.jungle.lsma.charging.pricing;</pre>
10	<pre>import com.sun.java.swing,JComponent; /**This establishes the operations used for communication between the customer system and the downloaded algorithm. @author Mike Rizzo*/ public interface Tariff { Jcomponenet getGUI(); Float getPrice (int pkin, int pkout, int cng);</pre>
15	

25792CON.DOC

Table 2

```
package algorithms.linear;

5      import com.sun.java.swing.JComponent;
      import com.sun.java.swing.JTextField;

      import com.bt.jungle.lsma.charging.pricing.Tariff ;

10     public class LinearAlgorithm implements Tariff {
          private float rate;
          private LinearAlgorithmDisplay display;
          public LinearAlgorithm ( ) {
15              display = new LinearAlgorithmDisplay ( );
              setRate (new Float(1));
          }
          public float getPrice (int pkin, int pkout, int cng) {
              return (pkin + pkout + cng) * rate;
          }
20     public JComponent getGUI( ) { return display; }
          public void setRate (Float f) {
              rate = f.floatValue( ) ;
              display.setRate (rate) ;
          }
25     }
```

Table 3

```

// Generated by Together

5    package algorithms.linear;

    import com.sun.java.swing.JPanel;
    import com.sun.java.swing.JTextField;
    import com.sun.java.swing.Box;
10   import com.sun.java.swing.JLabel;

    public class LinearAlgorithmDisplay extends JPanel {
        private JTextField tfRate = new JTextField (4);
        public LinearAlgorithmDisplay ( ) {

15             Box vbox = Box.createVerticalBox ( ) ;

            Box hbox = Box.createHorizontalBox ( ) ;
            hbox.add (Box.createHorizontalGlue ( ) ) ;
20             hbox.add (new JLabel ("Rate:"));
            hbox.add (Box.createHorizontalGlue( ));
            hbox.add (tfRate);
            tfRate.setEditable (false);
            hbox.add (Box.createHorizontalGlue( ))
25             vbox.add (hbox);

            add (vbox());
        }
        public void setRate (float f) {
30             tfRate.setText (String.valueOf (f));
        }
    }

```

Table 4

```

// Generated by Together
package algorithms.linear;

import com.sun.java.swing.*;
import java.awt.event.*;

import com.bt.jungle.lsma.charging.pricing.provider.*;
import com.bt.jungle.util.*;

public class LinearAlgorithmGUI extends JPanel {
    private JTextField tfRate = new JTextField ( );
    private TuningMessageListener tuningMessageListener;
    private final static String DEFAULT_RATE = "1.0";
    public LinearAlgorithmGUI (TuningMessageListener tml) {
        tuningMessageListener = tml;
        tfRate.setText (DEFAULT_RATE);

        Box vbox = Box.createVerticalBox( );

        Box hbox = Box.createHorizontalBox ( );
        hbox.add (new JLabel ("Rate:"));
        hbox.add (Box.createHorizontalGlue( ));
        hbox.add (tfRate);
        hbox.add (Box.createHorizontalGlue( ));
        hbox.add (hbox);

        JButton bTransmit = new JButton ("Transmit");
        bTransmit.addActionListener (
            new ActionListener ( ) {
                public void actionPerformed (ActionEvent e) {
                    transmit ( ) ;
                }
            }
        );
        hbox = Box.createHorizontalBox( );
        hbox.add. (Box.createHorizontalGlue( ));
        hbox.add (bTransmit);
        hbox.add (Box.createHorizontalGlue( ));
        vbox.add (hbox);

        add (vbox);
    }
    void transmit ( ) {
        try {
            Float f = new Float (tfRate.getText( ));
            Object args [ ] = new Object [1];
            Args [0] = f;
            TuningMessageListener.notify(

```

```

        new Invocation ("SetRate", args)
    );
    }
    catch (Exception e) {
5      e.printStackTrace ( );
    }
}
}

```

Table 5

```

package algorithms.exp;

5      import com.sun.java.swing.JComponent;
      import com.sun.java.swing.JTextField;

      import com.sun.java.lsma.charging.pricing.Tariff;

10     public class ExpAlgorithm implements tariff {
          private float min;
          private float base;
          private float divisor;
          private ExpAlgorithmDisplay display( );
15     public ExpAlgorithm( ) {
          display = new ExpAlgorithmDisplay ( );
          setMin (new Float (1));
          setBase (new Float (2));
          setDivisor (new Float (10));
20     }
          public float getPrice (int pkin, int pkout, int cng) {
          return min + (float)math.pow (base,(pkin + pkout + cng)/divisor);
          }
          public JComponent getGUI( ) {return display; }
25     public void setMin (Float f) {
          min = f.floatValue( );
          display.setMin(min);
          }
          public void setBase (Float f) {
30     base = f.floatValue( );
          display.setBase(base);
          }
          public void set Divisor (Float f) {
          base = f.floatValue( );
          display.setBase (divisor);
35     }
      }

```


Table 7

```

// Generated by Together

5  package algorithms.exp;

import java.awt.GridLayout;
import com.sun.java.swing.*;
import java.awt.event.*;

10 import com.bt.jungle.lsma.charging.pricing.provider.*;
import com.bt.jungle.util.*;

public class ExpAlgorithmGUI extends JPanel {
15     private JTextField tfMin = new JTextField ( );
    private JTextField tfBase = new JtextFfield ( );
    private JTextField tfDivisor = new JTextField ( );

    private TuningMessageListener tuningMessageListener;
20     private final static String DEFAULT_MIN = "1.0";
    private final static String DEFAULT_BASE = "2.0";
    private final static String DEFAULT_DIV = "10.0";

    public ExpAlgorithmGUI (TuningMessageListener tml) {
25         tuningMessageListener = tml;
        tfMin.setText (DEFAULT_MIN);
        tfBase.setText (DEFAULT_BASE);
        tfDivisor.setText (DEFAULT_DIV);

30         Box vbox = Box.createVerticalBox ( );

        vbox.add (new JLabel ("price = min + pow (base,
(pkIn + pkout + cng)/divisor)"));

35         vbox.add (Box.createVerticalGlue ( ));

        JPanel panel = new JPanel (new GridLayout (3,2));
        panel.add (new JLabel ("Minimum"));
        panel.add (tfMin);
40         tfMin.addActionListener (
            new ActionListener ( ) {
                public void actionPerformed (ActionEvent e) {
                    transmit ("setMin", tfMin);
                }
            }
45         );
        panel.add (new JLabel ("Base"));
        panel.add (tfBase);
        tfBase.addActionListener (
50         new ActionListener ( ) {

```

00674720-4560


```

    public void actionPerformed (ActionEvent e) {
        transmit ("setBase", tfBase);
    }
}
};
5   panel.add (new JLabel ("Divisor"));
    panel.add (tfDivisor);
    tfDivisor.addActionListener (
        new ActionListener ( ) {
10       public void actionPerformed (ActionEvent e) {
            transmit ("setDivisor", tfDivisor);
        }
    }
);
15   vbox.add (panel);

    add (vbox)
}
void transmit (String m, JTextField tf) {
20     try {
        Float f = new Float (tf.getText ( ));
        Object args [ ] = new Object [1];
        args [0] = f;
        tuningMessageListener.notify(
25         new Invocation (m, args)
        );
    }
    catch (Exception e) {
30         e.printStackTrace ( );
    }
}
}

```

00574730 140500

14 07 07 00

31

CLAIMS

1. A method of operating a communications network, including automatically varying at a customer terminal, depending on network loading as detected at the customer terminal, a tariff for network usage by the customer terminal.
2. A method according to claim 1, including detecting at the customer terminal a network performance parameter which depends on network loading, and varying the tariff depending on the network performance parameter.
3. A method according to claim 2, in which the network is a packet network and the network performance parameter is the number of packets lost in transmission between a data source and the customer terminal.
4. A method according to claim 1, including detecting a congestion signal at the customer terminal and varying the tariff in response to the congestion signal.
5. A method according to claim 4, including reading a congestion signal at the customer terminal from a data packet received at the customer terminal.
6. A method according to claim 4 or 5, including generating a congestion signal at a router in the network in response to the detection of congestion at the router.
7. A method according to any one of the preceding claims, including making a first relatively smaller increase in the tariff when congestion is first detected, and making at least one further, relatively larger increase, if the congestion persists
8. A method according to any one of the preceding claims, including programming a decision agent at the customer terminal with user-determined price criteria, and comparing a price calculated using the tariff with the said price criteria.
9. A method according to any one of the preceding claims, including distributing a tariff algorithm via the communications network to a plurality of terminals and

14 07 07 00

32

calculating at each terminal, using the tariff, a charge for network usage by the terminal.

10. A method according to claim 9, further comprising steps, carried out by the
5 network operator, of:

intermittently sampling traffic between a customer terminal and the network,
and as part of the sampling recording network loading affecting the customer
terminal; and

for the sampled traffic comparing a charge calculated by the customer
10 terminal and an expected charge and detecting thereby any discrepancy.

11. A method according to any one of the preceding claims, in which when the
customer terminal detects congestion in data transmitted to the customer terminal
from a data source via the network, the customer terminal returns a congestion
15 notification signal to the data source.

12. A method according to any one of the preceding claims, including at a customer
terminal, selecting a period of time for which the tariff is to be fixed and paying a
premium depending on the duration of the said period.

20

13. A method of operating a communications network including applying to
customer terminals a tariff for network usage, varying the tariff with time; at a
customer terminal, selecting a period of time for which the tariff is to be fixed; and
paying a premium depending on the duration of the said period.

25

14. A communications network including :

means for detecting network loading locally at a customer terminal; and

means responsive to the said means for detecting arranged automatically to
vary a tariff for network usage by the customer terminal.

30

15. A customer terminal for use in a communications network, the customer
terminal including:

M 07.07.00

33

means for detecting loading of a network to which, in use, the customer terminal is connected;

means responsive to the said means for detecting and arranged automatically to vary a tariff for network usage by the customer terminal.

5

14. A customer terminal for use in a communications network, the customer terminal including one or more processors arranged to carry out the following steps in sequence:

detecting loading of resources in a network to which the customer terminal is
10 connected; and

automatically varying in dependence on the detected loading a tariff for network usage by the customer terminal.

15. A method according to any one of claims 1 to 13, in which the tariff is
15 varied only if the terminal fails to reduce its output in response to detected congestion.

AMENDED SHEET

ABSTRACT

In a communications network, loading of network resources is detected locally at a customer terminal, and a tariff for network usage is varied automatically depending
5 on the usage.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)	(31)	(32)	(33)	(34)	(35)	(36)	(37)	(38)	(39)	(40)	(41)	(42)	(43)	(44)	(45)	(46)	(47)	(48)	(49)	(50)	(51)	(52)	(53)	(54)	(55)	(56)	(57)	(58)	(59)	(60)	(61)	(62)	(63)	(64)	(65)	(66)	(67)	(68)	(69)	(70)	(71)	(72)	(73)	(74)	(75)	(76)	(77)	(78)	(79)	(80)	(81)	(82)	(83)	(84)	(85)	(86)	(87)	(88)	(89)	(90)	(91)	(92)	(93)	(94)	(95)	(96)	(97)	(98)	(99)	(100)
1	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100




Figure 1

2/2

Figure 2a




Figure 2b





Figure 2c



00901F 02442560


Figure 3

0	1	2	3	4	5	6	7
IN			PHB				ECN

3/8

09/674720

Fig.4a.



5/7

09/674720

009011 0247960

Figure 4b




Fig.5a.

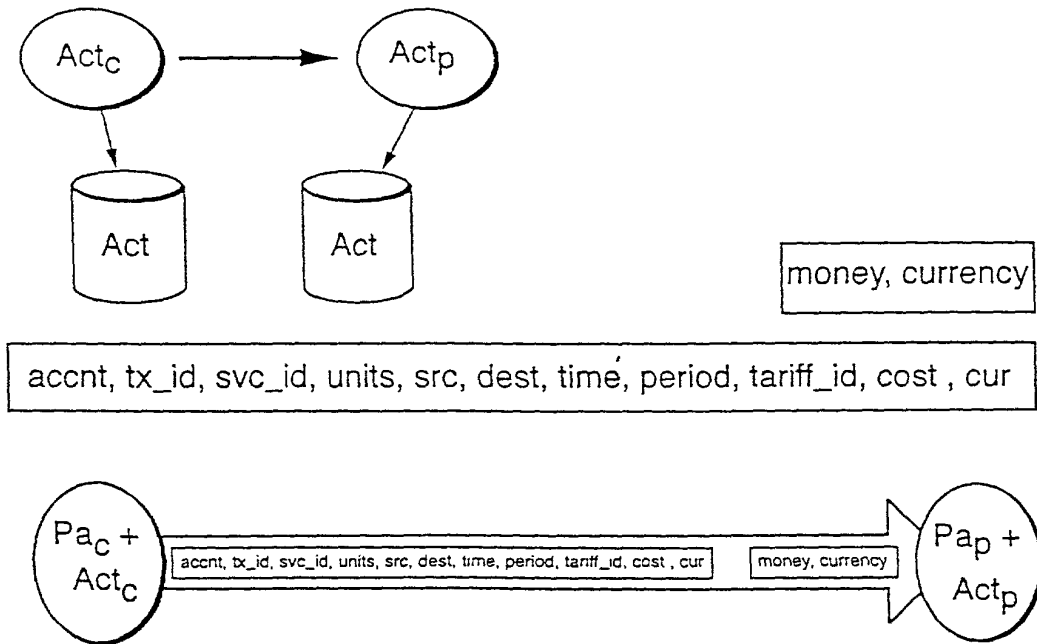


Fig.5b.

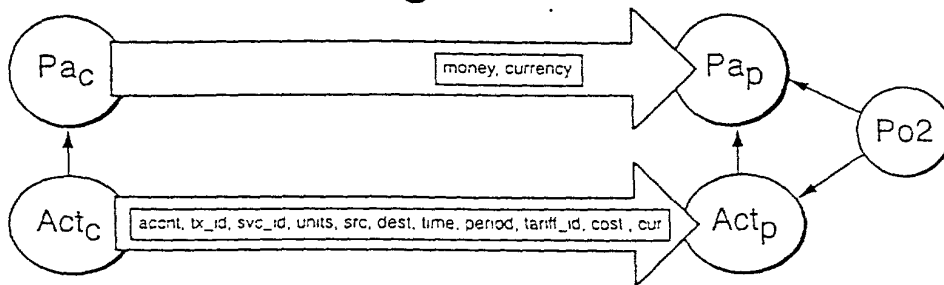
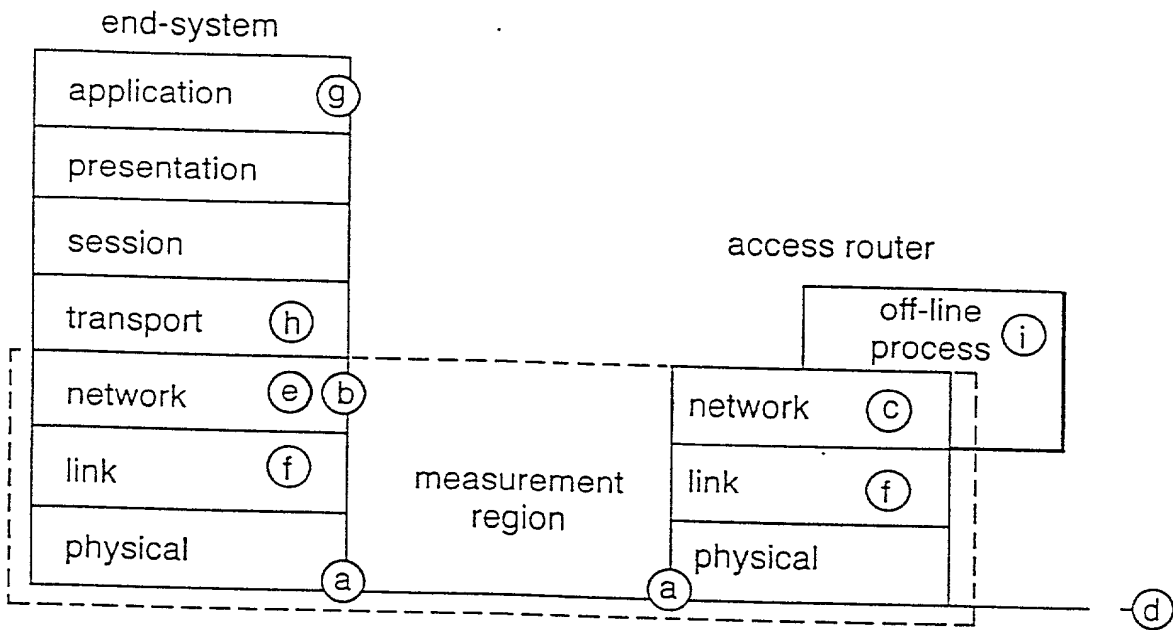


Fig.6.



09/674720

Figure 7




Figure 3a

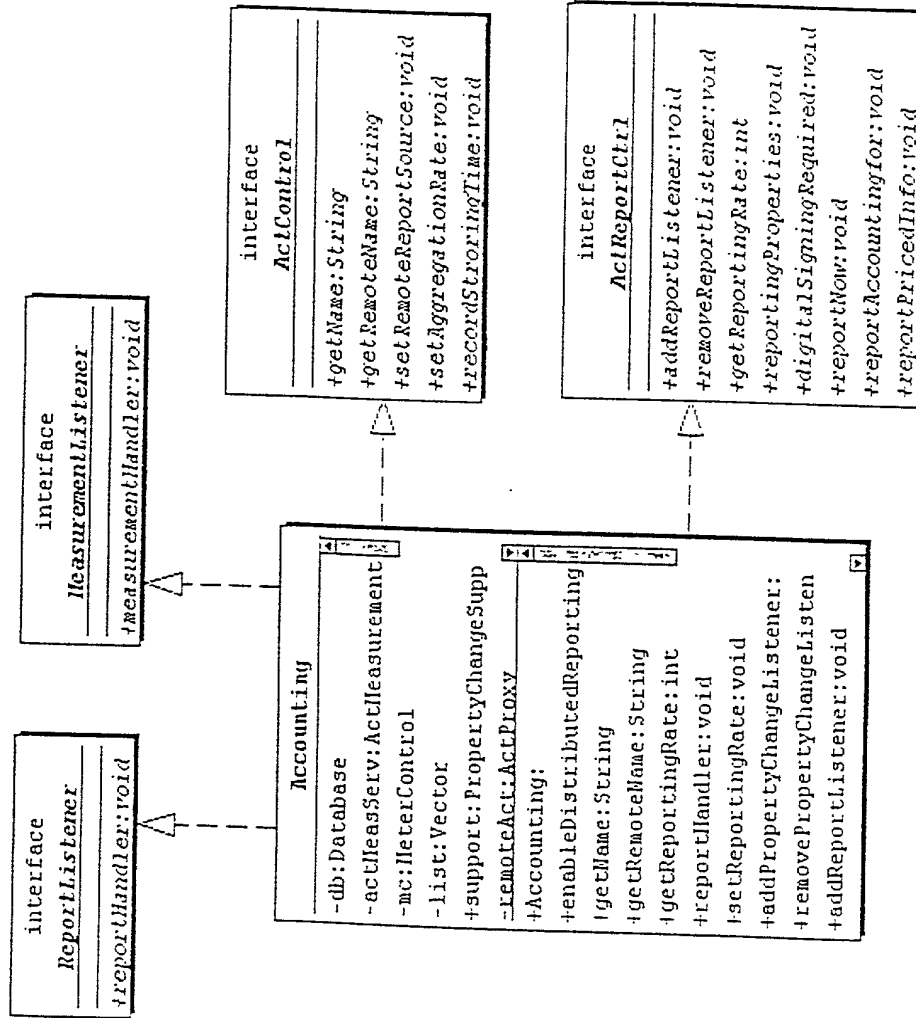


Figure 8b

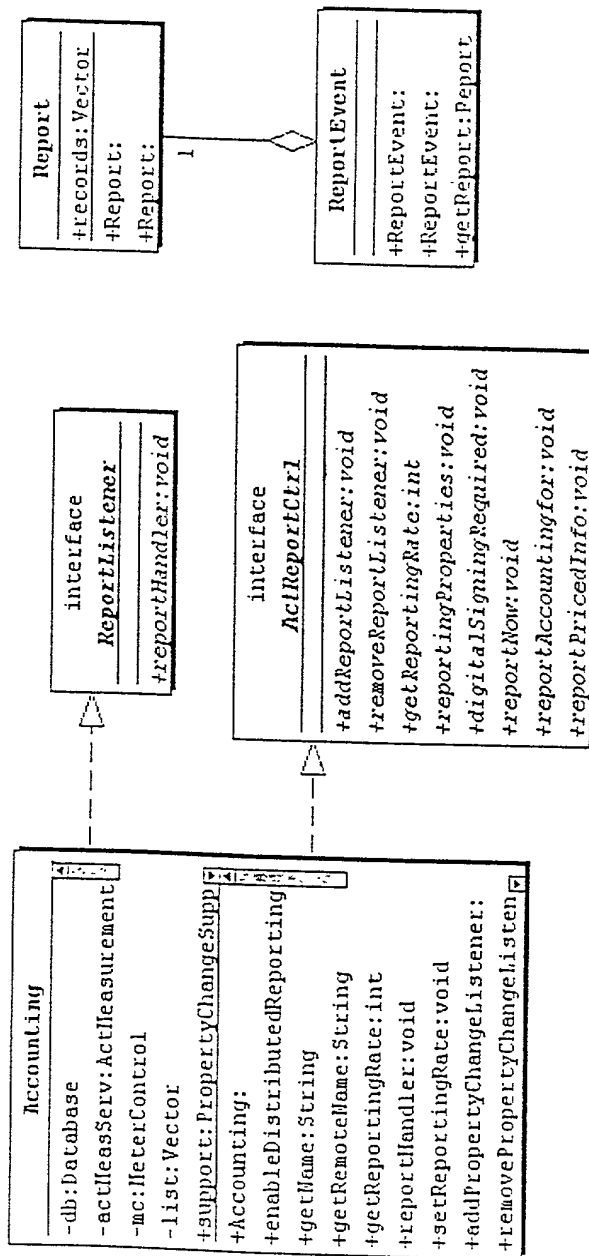


Figure 69

BTInternet Internet Accounting Control Platform

Local Platform ID	BTInternet
Local Meter ID	local
Local Reporting Rate	1000

Reporting Source: Add Remove

☐ Report priced data
☐ Security
☐ Encryption required
☐ Digital signing Required

Update

09/674720

Figure 8c

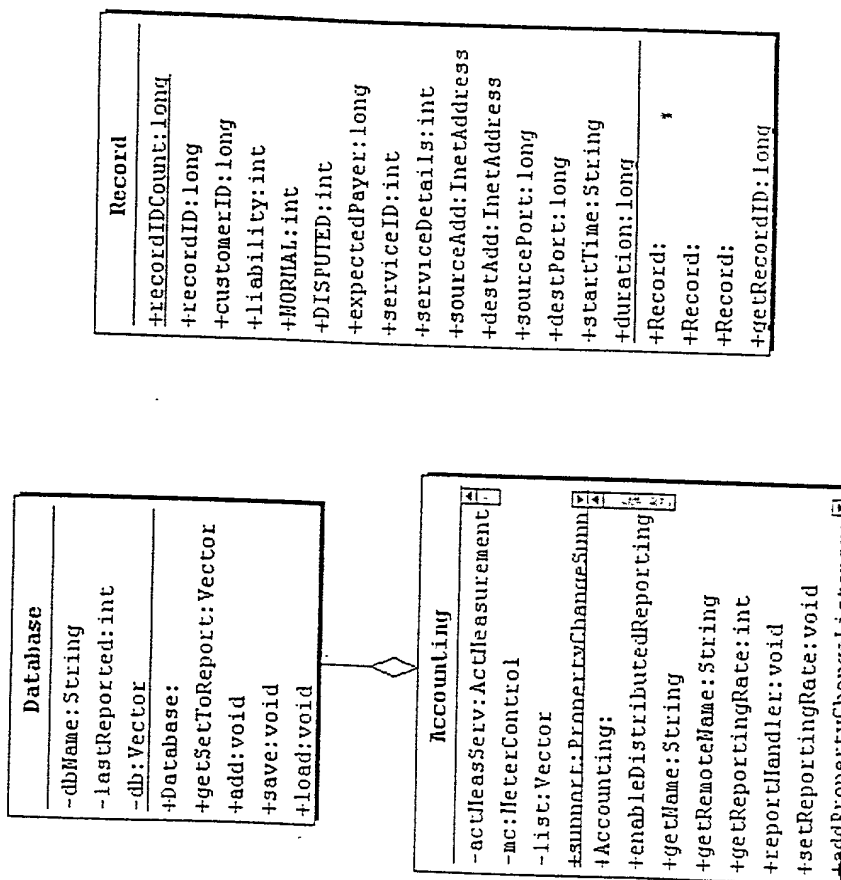


Figure 3d

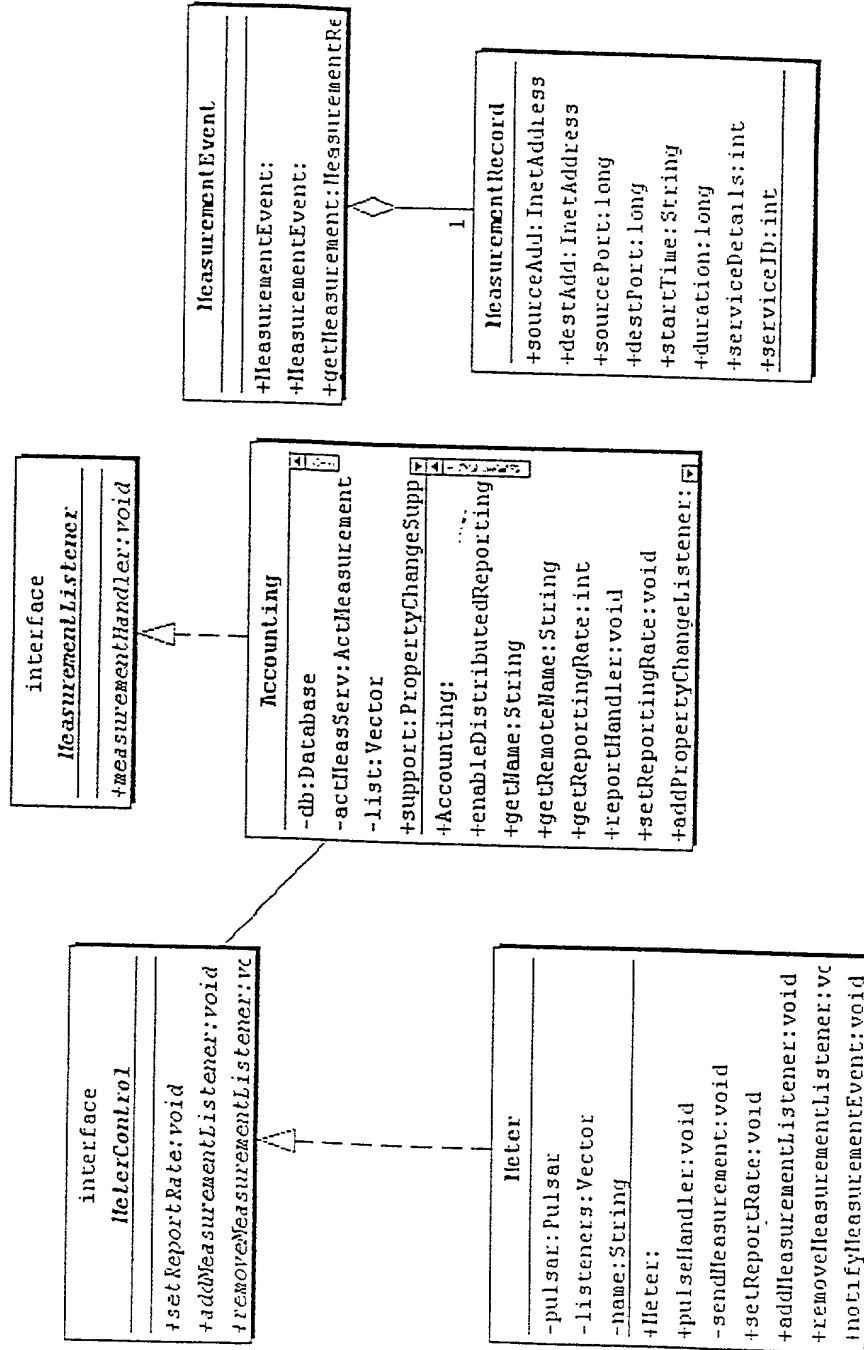
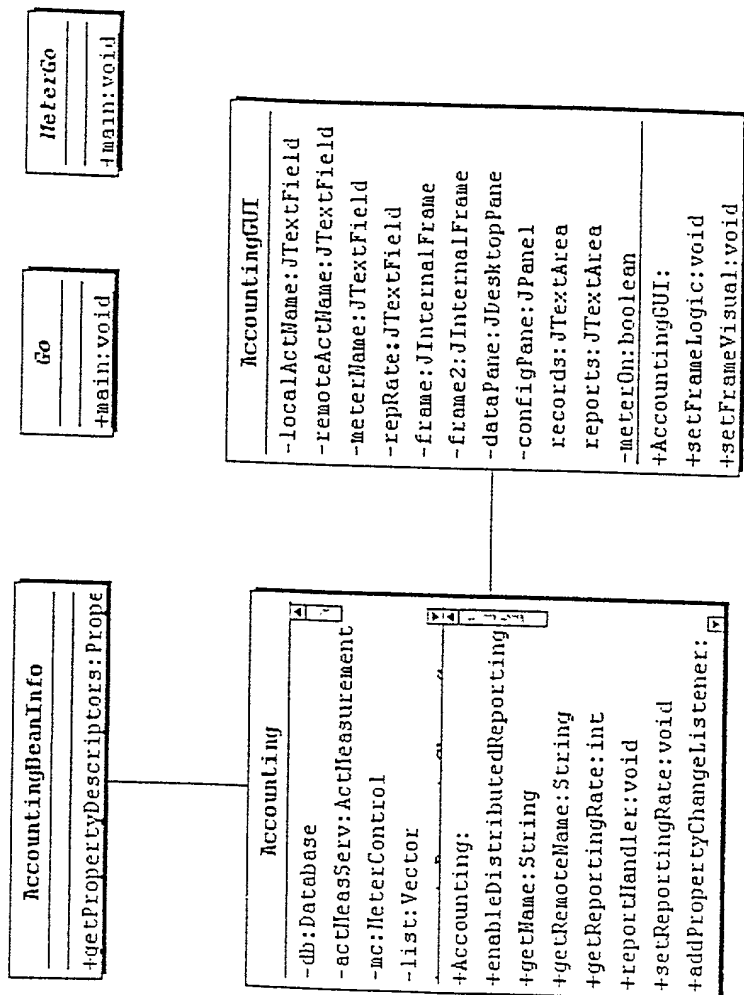


Figure 3e



File
(25792)





Figure 11



09/07/20

RULE 63 (37 C.F.R. 1.63)
DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

As a below named inventor, I hereby declare that my residence, post office address and citizenship are as stated below next to my name, and I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **COMMUNICATIONS NETWORK** the specification of which (check applicable box(es)):

☐ is attached hereto.

☐ was filed on _____ as U.S. Application Serial No. _____

☒ was filed as PCT international application No. PCT/ GB99/01773 on 04 June 1999

and (if applicable to U.S. or PCT application) was amended on _____

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to the examination of this application in accordance with 37 C.F.R. 1.56. I hereby claim foreign priority benefits under 35 U.S.C. 119/365 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed or, if no priority is claimed, before the filing date of this application:

Prior Foreign Application(s): Application Number	Country	Day/Month/Year Filed
9812161.9	GREAT BRITAIN	05 JUNE 1998
98309609.0	EUROPE	24 NOVEMBER 1998
9825723.1	GREAT BRITAIN	24 NOVEMBER 1998
9902052.1	GREAT BRITAIN	29 JANUARY 1999
9902648.6	GREAT BRITAIN	05 FEBRUARY 1999

I hereby claim the benefit under 35 U.S.C. 119(e) of any United States provisional application(s) listed below.

Application Number	Day/Month/Year Filed

I hereby claim the benefit under 35 U.S.C. 120/365 of all prior United States and PCT international applications listed above or below and, insofar as the subject matter of each of the claims of this application is not disclosed in such prior applications in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. 1.56 which occurred between the filing date of the prior applications and the national or PCT international filing date of this application:

Prior U.S./PCT Application(s): Application Serial No.	Day/Month/Year Filed	Status: patented, pending, abandoned
PCT/GB99/01773	04 June 1999	PENDING

I hereby declare that all statements made herein of my own knowledge are true and that statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon. And I hereby appoint NIXON & VANDERHYTE P.C., 1100 North Glebe Road, 8th Floor, Arlington, VA 22201-4714, telephone number (703) 816-4000 (to whom all communications are to be directed), and the following attorneys thereof (of the same address) individually and collectively my attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith and with the resulting patent: Arthur R Crawford, 25327; Larry S. Nixon, 25640; Robert A. Vanderhye, 27076; James T. Hosmer, 30184; Robert W. Faris, 31352; Richard G. Besha, 22770; Mark E. Nusbaum, 32348; Michael J. Keenan, 32106; Bryan H. Davidson, 30251; Stanley C. Spooner, 27393; Leonard C. Mitchard, 29009; Duane M. Byers, 33363; Paul J. Henon, 33626; Jeffry H. Nelson, 30481; John R. Lastova, 33149; H. Warren Burnam Jr., 29366; Thomas E. Byrne, 32205; Mary J. Wilson, 32955; J. Scott Davidson, 33489; Alan M. Kagen, 36178; William J. Griffith, 31260; Robert A. Molan, 29834.

1. Inventor's Signature: Robert J. Briscoe Date: 9 JUL 1999
 Inventor: ROBERT J BRISCOE BRITISH
 (first) (MI) (last) (citizenship)
 Residence: (city) SUFFOLK (state/country) GREAT BRITAIN
 Post Office Address: HOME FARM, PARHAM, WOODBRIDGE, SUFFOLK, GREAT BRITAIN
 (Zip code) IP13 9NW

2. Inventor's Signature: Michael J. Rizzo Date: 6/7/99
 Inventor: MICHAEL J RIZZO BRITISH MALTESE
 (first) (MI) (last) (citizenship)
 Residence: (city) SUFFOLK (state/country) GREAT BRITAIN
 Post Office Address: 12 DEWAR LANE, KESGRAVE, IPSWICH, SUFFOLK, GREAT BRITAIN
 (Zip code) IP5 2GJ

FOR ADDITIONAL INVENTORS, check box ☐ and attach sheet with the same information and signature and date for each.